Exploiting System z Innovation for Mainframe-based MFT with IBM Sterling Connect:Direct for z/OS

Session : 15674

Speaker: Mike Watley
Advisory Software Engineer
IBM Software Group
mewatley@us.ibm.com

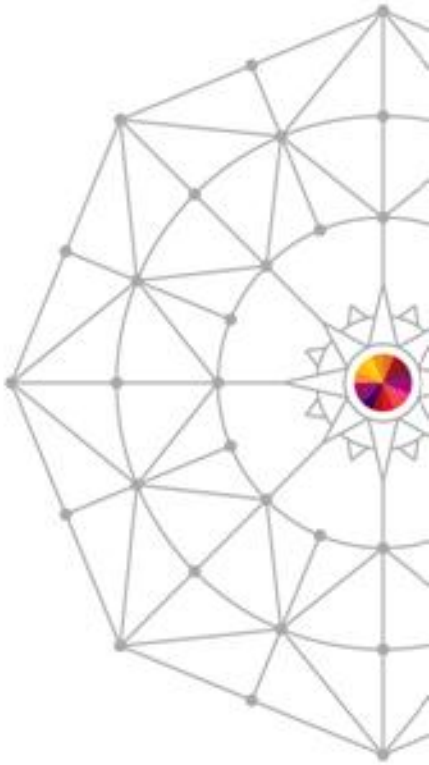

#SHAREorg

SHARE in Pittsburgh 2014

Disclaimer

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Exploiting System z Innovation

- Agenda
  - IBM System Storage DS8000 series and zFBA feature
  - IBM zEnterprise Data Compression (zEDC) for z/OS and zEDC Express Accelerator
  - Security and Compliance Updates
  - Questions and Answers

Data growth creates cross-platform file transfer challenges

Customer Challenges	Current technologies	Sterling Connect:Direct with zFBA
• Customers require large (e.g., petabytes) mission-critical Operational Data Store environments. • Data from z/OS needs be copied from the ODS frequently (peak data copy can exceed one Terabyte per hour) to large UNIX servers. • A TCP/IP solution, based on a dedicated IP network, or storage-based data sharing solution, cannot handle the volumes, causing significant server load on the both System z and UNIX servers. • Aggressive Service Level Objectives (SLOs) stipulate specific time windows for copying this data. Failure to meet these windows can adversely affect the total business process.	• Enterprises use their TCP/IP network to transmit bulk data between open systems server(s) and System z, which can create TCP/IP network gridlock. • This gridlock can disturb necessary business processes needed to maintain and continue growth. **Problems this creates:** • Elapsed time to move the data is becoming a critical bottleneck • CPU Utilization can be high enough to cause disruptions to other workloads.	**A better approach:** IBM Sterling Connect:Direct for z/OS and zFBA on IBM DS8000 storage solutions • Reduces the stress off TCP/IP network by moving data at the channel level • Reduce latency required to handle data volume • Reduce CPU utilization for file transfers, which could be as high as 50 percent. **Benefits:** • Fast ROI • Reduced file copy elapsed time • Offloading the file transfer processing from the host saves critical cycles • Reduced monitoring staff due to built in functionality.

Exploiting the zFBA

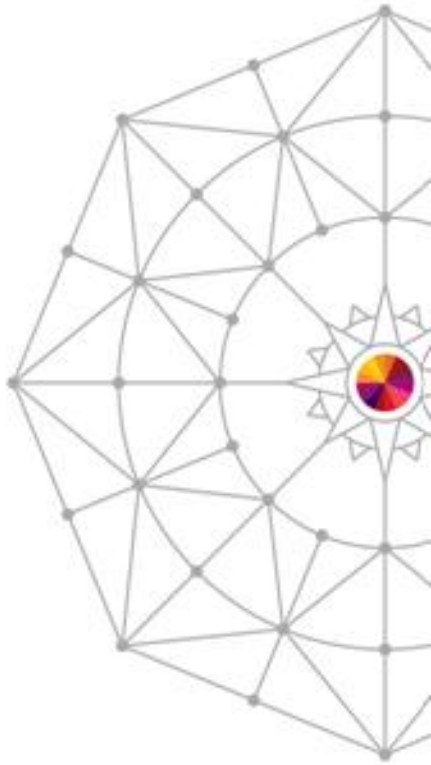

#SHAREorg

Exploiting the zFBA devices

- C:D utilizes the channel connected DS8000 series storage device for high speed data transfer between z/OS and Distributed UNIX
- Exploits the z/OS Distributed Data Backup (zDDB) Feature that allows definition of zFBA devices
- Communication with remote C:D via TCP while data transfer via the zFBA devices
- C:D allocates a number of zFBA devices using a new IOS API, IOSFBA
- Allows data to be written to one device while data is consumed from the other
  - C:D switches back and forth between allocated devices until all data has been transferred and consumed
- Data on the zFBA device is short lived, available for the remote C:D to consume then devices are cleared before being unallocated
- C:D z/OS defines a range of zFBA devices in the Network Configuration Map or NETMAP and controls allocation of the zFBA devices regardless of which end-point initiated the Process.

Offloading file transfer from TCP/IP network

A sample communication flow with zFBA

1. C:D z/OS establishes the TCP/IP connection with its remote partner as its normal communication path.
2. C:D z/OS utilizes the channel connected DS8800 as its data path.
3. C:D z/OS utilizes an IOS API to allocate FBA devices on the DS8800
4. C:D z/OS writes/reads I/O buffer to these allocated devices.
5. C:D z/OS communicates via the IP connection to the remote C:D node that data is available on said device.
6. C:D UNIX processes the device as a SCSI device to read/write the data then write the output to the destination file.

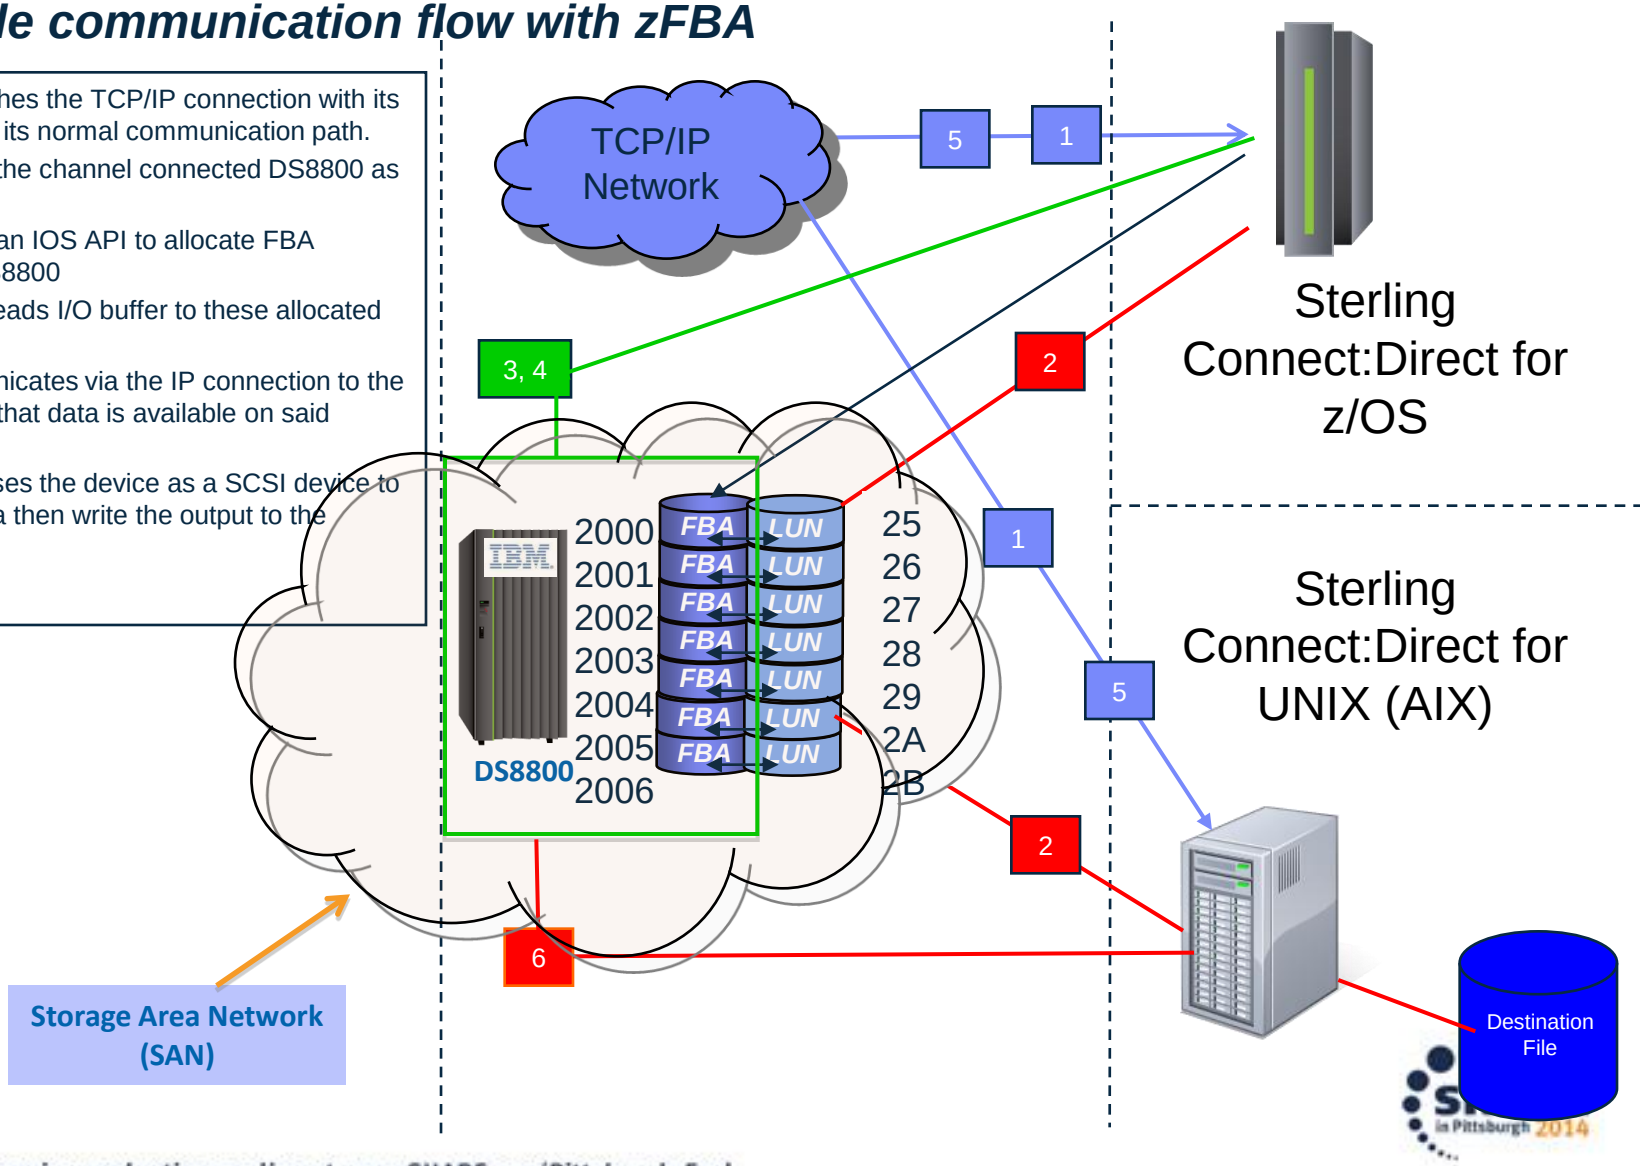

Support for zFBA devices

- Transfer between C:D nodes
  - C:D z/OS to/from C:D UNIX (AIX)
  - C:D z/OS to/from C:D z/OS
    - C:D z/OS nodes must be running in different LPAR(s)
  - C:D Process parameters must request the zFBA feature
- Supports all C:D feature functionality
  - Statistics logging
  - Compression
  - Automatic restart of session or allocation failure
  - Checkpoint / Restart
  - All supported File types
  - All supported record formats
- Current exceptions
  - Secure Plus
  - PDS / PDSE

Hardware Dependencies

- Storage devices must conform to the IBM Fixed Block Architecture (FBA) command set as implemented by the IBM DS8000 family with the z/OS Distributed Data Backup (zDDB) Feature
- zDDB Feature is a no charge microcode feature for IBM DS8000 family
- For DS/8700 and DS/8800 the updates/fixes are available in Release 6.3SP5 or higher.
  - For DS8700, MCL 76.31.63.0 or higher
  - For DS8800, MCL 86.31.78.0 or higher

- For DS/8870 updates/fixes are available in Release 7.1.5 or higher
  - For DS8870, MCL 87.10.97.0 or higher

z/OS Software Dependencies

- Available support via z/OS V2R1
- Required PTFs and APARs for support
  - UA70715 for z/OS 2.1 (HBB7790)
  - UA70714 for z/OS 1.13 (HBB7780)
  - UA70716 for z/OS 1.13.1 (JBB778H)
  - APAR OA44034 for z/OS 1.13 (HBB7780) and z/OS 2.1 (HBB7790)

Steps for Implementation

- Configuration of zFBA devices
  - Determine number of devices required
    - Device pairs determines the number of concurrent transfers
  - Minimum size for each zFBA device
    - 32MB minimum
    - Larger devices is a waste
  - Configure the hardware

- Define zFBA device range in the C:D NETMAP Configuration
  - z/OS to z/OS, both nodes require the same ZFBA keyword and range
  - C:D UNIX NETMAP does not define the device range

```
ADJACENT.NODE=((CD.UNIX,4399,10.20.112.2,TCP)   -
   ENVIRONMENT=UNIX  -
   PARSESS=(25 1)   -
   ZFBA=(2000,2006)   )
```

Steps for Implementation - Continued

- Initialize C:D with zFBA support
  - Specify ZFBA Initialization Parameter for both z/OS and UNIX
    - ZFBA = YES
- Define ZFBA keyword in the C:D Process

```
ZFBA01  PROCESS SNODE=CD.UNIX
PUSH    COPY FROM (PNODE DSN= MWATL1.F80.DATA1111  -
                   DISP=(SHR) )    -
                TO   (SNODE DSN= /sci/users/mwatl1/data/data1.txt  -
                   DISP=(RPL) )     -
              CKPT=30M COMPRESS EXTENDED ZFBA=2
```

zFBA Performance

Protocol	Transfer Direction	File Size	Transfer Rate (MBPS)	Elapse Time (Seconds)	CPU (seconds)
zFBA	Push to AIX	8GB	68.83	116	8.190
TCP	Push to AIX	8GB	58.28	137	10.545
zFBA	Push (Stripped)	8GB	156.55	51	8.123
TCP	Push (Stripped)	8GB	63.87	125	13.012

Explosive growth in data
Every day 2.5 quintillion bytes of data are created

Data needs to be shared cross platform

- Today's modular business applications are deployed on heterogeneous platforms
- Applications need to be able to efficiently share and process large files

Managing and storing data can push expenses up for IT

- Many compression applications consume CPU resources
- Batch workloads and some data types can drive up compression requirements
- Not being able to compress files can contribute to added disk expense

Many types of compression are used today

- Many offerings for archive – creating data that will be infrequently accessed
- Some is done to create smaller files allowing for faster network transfers
- Industry standard compression offerings in market like zlib
- IBM System z® offers hardware compression on all System z processor chips

…. Still massive amounts of data are not being compressed and need a different approach to compress more "active" data

IBM zEnterprise Data Compression (zEDC) New data compression offering that can reduce resource usage

NEW

What is it?

- ✓ *An combined software (z/OS V2.1) and hardware (zEDC Express) solution designed to help reduce resource consumption, disk utilization and optimize cross platform exchange of data*

How is it different

- **Performance:** Efficient alternative for larger files
- **Efficient:** Optimized algorithms scan text to locate the re-use of phrases and refers back to earlier references
- **Compatible:** Compatible with open zlib based compression – used today by Java™ and other applications including IBM Sterling Connect:Direct
- **Economical:** Savings expected in CPU resources, disk storage and network resources

Exploiting zEDC

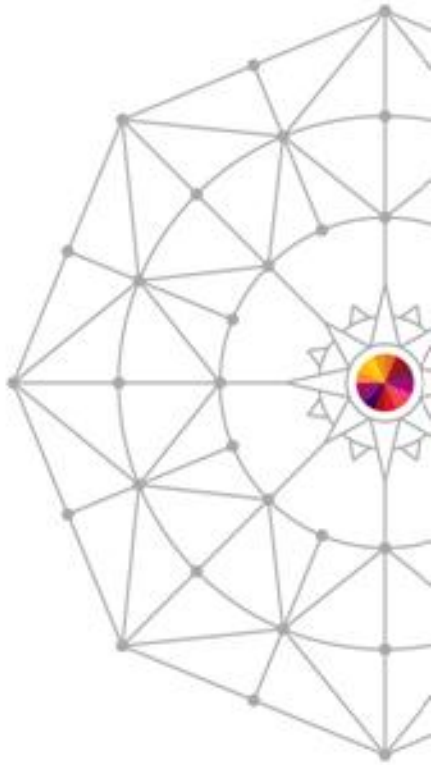

#SHAREorg

Exploiting zEDC Express Accelerator

- Support to exploit zEDC (Enterprise Data Compression) Express accelerator available for EC12/BC12 Servers
  - Requires z/OS V2R1 with APAR OA43256
  - Implements new version of ZLIB, v1.2.7
- Can greatly reduce the amount of CPU required to compress data
  - Like all Compression, data dependent
  - Support global control and Process Step level control
  - Data size must be at least 4096 (size threshold)
- RACF Access control required
  - FACILITY Class - FPZ.ACCELERATOR.COMPRESSION
  - ZILB INIT function performs verification on the hardware and RACF access control
- If hardware is available, ZLIB INIT is successful and the compress payload meets size threshold then the compression function is performed using the accelerator. Otherwise software ZLIB is used.

zEDC Express Accelerator – Benefits

- Software ZLIB 1.2.7 performs much better than ZLIB 1.2.3
  - Reduced CPU over ZLIB 1.2.3
  - Faster than ZLIB 1.2.3
- zEDC Express Accelerator reduces CPU and Improved Process thru-put
- Provides multiple compression options
  - ZLIB software compression
  - zIIP exploitation of ZLIB software compression
    - zEDC can not exploit the zIIP
  - zEDC Express Accelerator
- MVS Programming: Callable Services for High-Level Languages , SA23-1377-00
  - Overview , Planning and Requirements

zEDC Express Accelerator – Performance

			ZLIB Software			zEDC Hardware		
Data Type	Record Format	File Size	Comp %	CPU sec	Elapse time	Comp %	CPU sec	Elapse time
Text	VB	1GB	73.2	26.107	26.485	76.2	0.699	5.769
Text	FB	1GB	84.5	11.706	11.900	77.5	0.689	6.569
Binary	FB	80MB	-0.5	10.052	10.206	-0.2	0.085	0.510

Security and Compliance

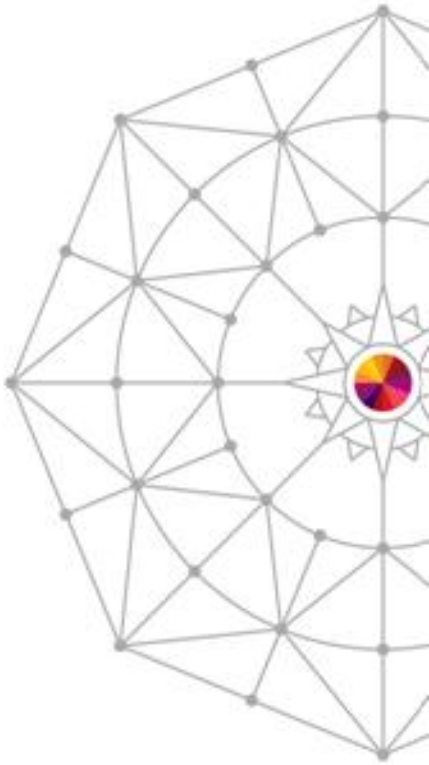

SecurePlus Security Compliance

- Updates of Secure Protocols
  - Add TLSv1.1 and TLSv1.2 Protocols
    - TLSv1.2 requires z/OS V2R1
    - PTF UA66870 and UA66872 for V1R13
  - S+ secure protocols now include SSLv3.0, TLSv1.0, TLSv1.1 and TLSv1.2
  - Removal of STS protocol
- Use of ICSF encryption services to protect PARMFILE contents
  - Removal of Certicom encryption toolkit
  - S+ Admin and C:D STC userid must be defined to CSFSERV Facility class
  - ICSF Random number service (CSNBRNG) requires Crypto-coprocessor
    - FMID HCR77A0 or higher, Crypto-coprocessor not required
- Implementation of Security Modes
  - FIPS Mode
  - NIST SP800-131a Transition Mode
  - NIST SP800-131a Strict Mode
  - NSA Suite B Profiles
    - 128 bit encryption cipher
    - 192 bit encryption cipher

SecurePlus Security Compliance

- Enforcement of Security Modes
  - C:D S+ enforces SP800-131a
    - Applied after successful Handshake
  - System SSL enforces NSA Suite B
    - Results in Failed Handshake

- Security Modes require FIPS
  - S+ Initialization will enable FIPS Mode when
    - FIPS = YES
    - FIPS = NO but Parmfile Enables FIPS
    - Refer to z/OS Cryptographic Service System Secure Socket Layer Programming, SC14-7495-00, for planning and implementation requirements for System SSL in FIPS mode

Exploiting System z Innovation for Mainframe-based MFT with IBM Sterling Connect:Direct for z/OS

Questions ?

Thank you!

#SHAREorg

Legal Disclaimer

-
- The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.
- References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.
- If the text contains performance statistics or references to benchmarks, insert the following language; otherwise delete:
  Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.
- If the text includes any customer examples, please confirm we have prior written approval from such customer and insert the following language; otherwise delete:
  All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.
- Please review text for proper trademark attribution of IBM products. At first use, each product name must be the full name and include appropriate trademark symbols (e.g., IBM Lotus® Sametime® Unyte™). Subsequent references can drop "IBM" but should include the proper branding (e.g., Lotus Sametime Gateway, or WebSphere Application Server). Please refer to http://www.ibm.com/legal/copytrade.shtml for guidance on which trademarks require the ® or ™ symbol. Do not use abbreviations for IBM product names in your presentation. All product names must be used as adjectives rather than nouns. Please list all of the trademarks that you use in your presentation as follows; delete any not included in your presentation. IBM, the IBM logo, Lotus, Lotus Notes, Notes, Domino, Quickr, Sametime, WebSphere, UC2, PartnerWorld and Lotusphere are trademarks of International Business Machines Corporation in the United States, other countries, or both. Unyte is a trademark of WebDialogs, Inc., in the United States, other countries, or both.
- If you reference Adobe® in the text, please mark the first use and include the following; otherwise delete:
  Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
- If you reference Java™ in the text, please mark the first use and include the following; otherwise delete:
  Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- If you reference Microsoft® and/or Windows® in the text, please mark the first use and include the following, as applicable; otherwise delete:
  Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.
- If you reference Intel® and/or any of the following Intel products in the text, please mark the first use and include those that you use as follows; otherwise delete:
  Intel, Intel Centrino, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- If you reference UNIX® in the text, please mark the first use and include the following; otherwise delete:
  UNIX is a registered trademark of The Open Group in the United States and other countries.
- If you reference Linux® in your presentation, please mark the first use and include the following; otherwise delete:
  Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.
- If the text/graphics include screenshots, no actual IBM employee names may be used (even your own), if your screenshots include fictitious company names (e.g., Renovations, Zeta Bank, Acme) please update and insert the following; otherwise delete: All references to [insert fictitious company name] refer to a fictitious company and are used for illustration purposes only.